

# DRAMScope: Uncovering DRAM Microarchitecture and Characteristics by Issuing Memory Commands

Hwiyong Nam<sup>†</sup>, Seungmin Baek<sup>†</sup>, Minbok Wi<sup>†</sup>, Michael Jaemin Kim<sup>†</sup>, Jaehyun Park<sup>†</sup>, Chihun Song<sup>‡</sup>,  
Nam Sung Kim<sup>‡</sup>, Jung Ho Ahn<sup>†</sup>

<sup>†</sup>Seoul National University, <sup>‡</sup>University of Illinois at Urbana-Champaign

<sup>†</sup>{hny4916, qortmdalss, homakaka, michael604, wogus20002, gajh}@snu.ac.kr, <sup>‡</sup>{chihuns2, nskim}@illinois.edu

**Abstract**—The demand for precise information on DRAM microarchitectures and error characteristics has surged, driven by the need to explore processing in memory, enhance reliability, and mitigate security vulnerability. Nonetheless, DRAM manufacturers have disclosed only a limited amount of information, making it difficult to find specific information on their DRAM microarchitectures. This paper addresses this gap by presenting more rigorous findings on the microarchitectures of commodity DRAM chips and their impacts on the characteristics of activate-induced bitflips (AIBs), such as RowHammer and RowPress. The previous studies have also attempted to understand the DRAM microarchitectures and associated behaviors, but we have found some of their results to be misled by inaccurate address mapping and internal data swizzling, or lack of a deeper understanding of the modern DRAM cell structure. For accurate and efficient reverse-engineering, we use three tools: AIBs, retention time test, and RowCopy, which can be cross-validated. With these three tools, we first take a macroscopic view of modern DRAM chips to uncover the size, structure, and operation of their subarrays, memory array tiles (MATs), and rows. Then, we analyze AIB characteristics based on the microscopic view of the DRAM microarchitecture, such as 6F<sup>2</sup> cell layout, through which we rectify misunderstandings regarding AIBs and discover a new data pattern that accelerates AIBs. Lastly, based on our findings at both macroscopic and microscopic levels, we identify previously unknown AIB vulnerabilities and propose a simple yet effective protection solution.

## I. INTRODUCTION

A deep understanding of DRAM microarchitecture and error characteristics is more important than ever; processing in memory (PIM) spotlighted [10], [33], [34], [51], [63], soft/hard error rate exacerbated [1], [23], [67], and yet another activate-induced bitflip (AIB) vulnerability discovered [39]. For instance, constructing secure and efficient AIB protection solutions without an accurate understanding of DRAM error behaviors linked to specific aspects of a DRAM microarchitecture would be undoubtedly challenging. Likewise, a detailed knowledge of the DRAM microarchitecture is essential in exploring efficient PIM architectures. However, the DRAM microarchitecture has undergone decades of optimizations to improve not only the cell density or energy efficiency but also the manufacturing yield and cost. Such optimizations are manufacturer-specific and proprietary [61], significantly hindering efforts to uncover the true DRAM microarchitecture and error characteristics.

To fill this critical gap, a large body of prior work has exploited creative reverse-engineering methodologies. They have relied on scarcely disclosed knowledge or assumptions [4], [17], [18], [21], [43], [44] to uncover error characteristics [4], [25], [29], [36], [39], [50], undefined DRAM operations [62], [82], microarchitectural components transparent to memory controllers, such as AIB protection solutions [9], [13] or on-die ECC [54], [55], to list a few. Nonetheless, we have found a number of previous efforts to discover the DRAM microarchitecture are limited in scope, outdated, or even misleading due to an insufficient understanding of the modern DRAM 6F<sup>2</sup> cell structure (see Figure 2), complex mapping of CPU physical addresses to DRAM addresses, and swizzling of CPU data within DRAM.<sup>1</sup>

In this paper, we conduct a comprehensive study to better understand the DRAM microarchitecture (*macroscopic* level) and AIB characteristics (*microscopic* level) of modern DRAM chips, leveraging three different reverse-engineering techniques and our recent knowledge of the aforementioned address mapping and data swizzling. Without a thorough understanding of the address mapping and data swizzling, attempting to control DRAM chips can lead to inconsistencies between the user’s intended access and the physical access. Similarly, comprehending the 6F<sup>2</sup> cell structure and the physical distances between cells and intervening gate types is essential for obtaining clearer insights from reverse-engineering efforts. We uniquely exploit this interplay by utilizing DRAM errors to uncover the DRAM microarchitecture while simultaneously leveraging our recent microarchitectural knowledge to investigate error characteristics.

**Reliable and cross-validatable reverse-engineering techniques (§III):** To reverse-engineer the DRAM microarchitecture without intrusive measures such as physical probing [4], [5], we use three techniques using standard DRAM commands in a controlled FPGA-based environment. The three techniques are as follows: (1) causing AIBs such as RowHammer [29] and RowPress [39], (2) performing in-memory row copy operations (RowCopy) [10], [62], and (3) inducing data retention errors. Analyzing the results obtained from these three techniques provides us with not only the

<sup>1</sup>DRAM internal data swizzling occurs as data collected from the subarray is reorganized to get transferred to the CPU. See § IV.

accurate error characteristics but also the hidden details of the DRAM microarchitecture. Furthermore, we highlight the challenges posed by intricate address mapping and data swizzling schemes, including row address remapping at individual DRAM chips, row address inversion at the registered clock driver (RCD) chip, and data pin (DQ) twisting. Although such information is often publicly disclosed [19], [21], they are scattered across documents and can easily be omitted, leading to incorrect analysis.

#### Macroscopic DRAM microarchitectural analysis (§IV):

We conduct a *macroscopic* analysis that does not require knowledge of the  $6F^2$  cell structure to reverse-engineer the *data swizzling* and identify previously unreported structural **Observations** at the subarray, row, and memory array tile (MAT) levels. **(O1)** We reconstruct the DRAM chip internal data swizzling based on our observation that horizontally adjacent victim cells affect AIB, which we elaborate later in §V. We observe that data within a single read are reorganized and collected from multiple MATs. **(O2)** We also identify the MAT *width*, or the number of cells constituting a single row in a MAT. **(O3)** For certain DRAM chips, two separate rows specified in physical address are coupled and activated together by a single row command. **(O4)** For all tested DRAM modules, the number of rows in a single subarray (*height*) is not a power of two, and multiple subarray heights can coexist even in a single chip. Also, we recognize a clear trend of increase in subarray height over DRAM generations. **(O5)** Certain DRAM chips combine two subarrays at the physical edge to work in tandem (edge subarrays), which is deducible from the open bitline structure [16]. **(O6)** The bit error rate (BER) by AIB is lower in edge subarrays, possibly due to the dummy bitlines.<sup>2</sup>

**Microscopic DRAM error analysis (§V):** With our *microscopic* analysis that exploits our knowledge of the  $6F^2$  cell structure, we present the following observations. **(O7)** There exists an alternating pattern in RowPress vulnerability, which reverses when either the row parity (even/odd) changes or the aggressor direction (up/down) changes. This faithfully reflects the  $6F^2$  cell structure. **(O8)** RowHammer also exhibits a similar alternating pattern, which is reversed when row parity, aggressor direction, or the written value (0/1) changes. **(O9)** RowHammer occurs at both types of cell access transistors, *i.e.*, neighboring and passing gates (§II-B). **(O10)** A victim cell is only susceptible from RowHammer to only one type of the gate at a time, which is reversed when the written value changes. **(O11)** Given a particular victim cell, its horizontally adjacent four victim cells’ data affect its RowHammer vulnerability, which becomes strongest at a distance of two. **(O12)** Similar horizontal influence exists in the aggressor row, which becomes weakest at distance two. We also find that a newly-discovered adversarial data pattern **(O13)** decreases the activation count that triggers the first bitflips ( $H_{\text{cnt}}$ ) in a victim row by up to 81% and **(O14)** exacerbates the overall bit error rate (BER) of the victim row by up to  $1.69\times$ .

<sup>2</sup>The edge subarrays of an open bitline structure utilize only half the bitlines, leaving the other half as dummy bitlines.

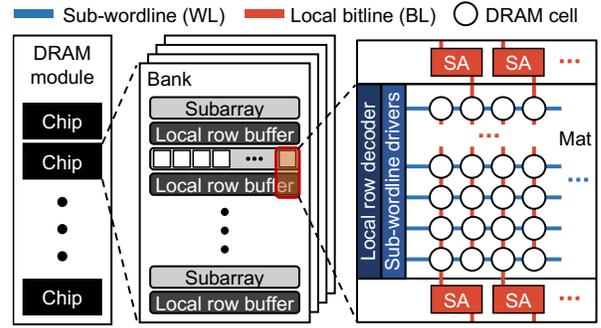


Figure 1. Conventional DRAM organization.

**New vulnerabilities and protection patches (§VI):** Based on the new observations at both macroscopic and microscopic levels, we identify previously unknown AIB vulnerabilities and propose a simple yet effective protection solution. First, we identify that edge subarrays and coupled rows can pose a new threat to existing AIB protection solutions, those driven by the memory controller. Besides, it can facilitate the attacker to succeed in a memory templating/massaging [30] phase, which is an essential part of a successful attack. Second, we demonstrate that, depending on the intention of the AIB attack, we can exploit an adversarial data pattern: (1) to achieve a specific cell’s bitflip or (2) to maximize the number of bitflips in the target row by collocating the row and column data/directional dependence. Lastly, countering this, we propose a simple yet effective data masking mechanism that can prevent the exploitation of such vulnerability.

## II. BACKGROUND

In this section, we present the background on DRAM organization, microarchitectures, and operations. We also overview DRAM activate-induced bitflips (AIBs).

### A. DRAM Organization

A DRAM module is hierarchically organized, from top to bottom, chips, banks, subarrays, MATs, and cells (see Figure 1). A DRAM cell consists of a capacitor and an access transistor, indexed by row/column address via the corresponding wordline (WL) and bitline (BL), respectively. A row decoder enables a specific WL, which turns on the access transistors that connect the cell capacitors to the sense amplifiers (SAs) via BLs. Each cell stores 1-bit data, and is classified into true-cell (anti-cell) if a charged state represents 1 (0) [36], [41], [77]. An SA senses and amplifies a voltage small difference in a pair of BLs and temporarily stores the value of a cell.

DRAM cells form a 2D array structure referred to as MAT, and an array of MATs constitute a single subarray. A single read/write command reads and writes data from one or more MATs. A subarray can have either an open or folded bitline structure, depending on whether a single SA is connected to both the upper and lower BLs (open) or not (folded) [16]. In an open BL structure, half BLs of a subarray share SAs with those of the *upper* subarray and the other half BLs of a subarray with those of the *lower* subarray.

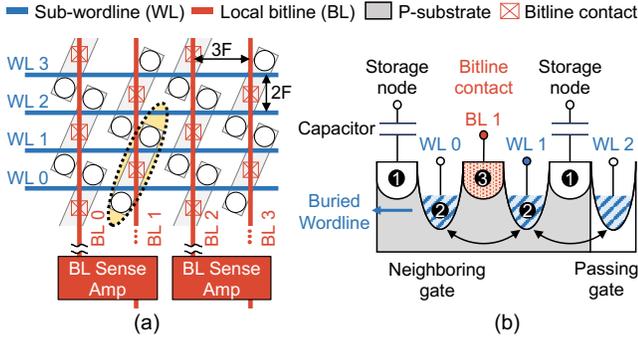


Figure 2. (a) DRAM  $6F^2$  cell structure and (b) the cross-section schematic of a saddle-fin transistor.

### B. $6F^2$ Cell Structure

Modern DRAM chips are primarily designed using a  $6F^2$  cell structure [15] for higher cell density, where  $F$  represents one-half of the minimum pitch (see Figure 2(a)). Figure 2(b) shows a cross-section view of the  $6F^2$  cell structure's highlighted region in Figure 2(a). The  $6F^2$  cell structure adopts a saddle-fin transistor, where a P-type substrate houses ① a pair of DRAM cells, each having a storage node and ② controlled by a buried WLs ( $WL0$  or  $WL1$ ). ③ These two cells share a BL connected through a bitline contact (BC). When  $WL1$  is enabled, we denote the adjacent WL ( $WL2$ ) (which does not share the P-substrate with  $WL1$ ) and the other adjacent WL ( $WL0$ ) as the passing gate and the neighboring gate, respectively [15], [52], [84].

### C. DRAM Operation

To access data, the memory controller (MC) sends an activate (ACT) command to enable a WL and connect the corresponding row of DRAM cells to BLs. The SAs and BLs are initially precharged to  $V_{dd}/2$ . When cells are connected to BLs, charge sharing occurs. This causes a small deviation in the voltage level of the BL, which is amplified to  $V_{dd}$  or 0 by SA. During the activation of a DRAM row, the global row decoder selects a subarray and the local row decoder selects and drives the corresponding row and WL. When DRAM receives a read (RD) or write (WR) command, the sensed or to-be-written data pass through the local and global I/O, equipped with temporary buffers on its path (e.g., global dataline SA).

$\tau_{RCD}$  is the minimum time between an ACT command to the RD/WR command. The RD/WR command reads/writes data from/to the sensed row in the SAs. After completing read or write operations, the host sends a precharge (PRE) command to disable the activated row's WL, disconnecting cells from the BLs. Prior to the precharge, the voltage level of the cell must be restored to  $V_{dd}$  or 0. The required time to issue the PRE command after the ACT command is  $\tau_{RAS}$ . After a PRE command is issued, the SAs and BLs require  $\tau_{RP}$  time to restore the BL voltage to  $V_{dd}/2$ .

### D. DRAM Activate-Induced Bitflips (AIBs)

AIBs are the representative DRAM read disturbance errors wherein activation disturbs cells in adjacent rows and flips the

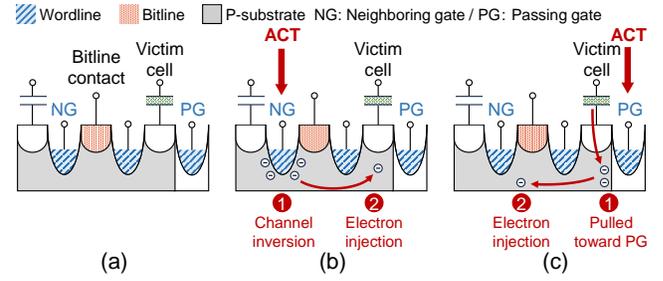


Figure 3. Mechanisms of activate-induced bitflips. PG and NG denote the passing gate and the neighboring gate, respectively.

states of the cells (see Figure 3). There are two mechanisms occurring AIBs: (1) electron migration (injection and capture) or (2) capacitive crosstalk [12], [58], [73], [86]. Based on bitflip mechanisms, the pair of an aggressor's WL and a victim cell is decided. The value of a victim cell can be flipped by electron migration from hammering the neighboring gate or by capacitive crosstalk from hammering the passing gate. Also, depending on the access pattern, AIBs can be caused by two attack patterns: (1) repetitive activation of the specific/aggressor row (RowHammer) or (2) activating the specific/aggressor row over an extended period (RowPress) [15], [29], [39].

The mechanism of AIBs depends on what type of gate the aggressor's WL is with the victim cell. Figure 3(b) demonstrates AIB by electron injection when the neighboring gate is turned on. When the aggressor's WL (the neighboring gate) is activated, ① electrons accumulate around the buried WL due to channel inversion. Upon deactivation, ② the accumulated electrons are spread out, while some are injected into the victim cell sharing the P-substrate. Figure 3(c) illustrates AIB by electron spreading originating from the passing gate. When the aggressor's WL (the passing gate) is activated, ① electrons are continuously attracted from the victim cell toward the passing gate. After the row is precharged, ② the electrons are spread out and some are injected into the active region, instead of returning to the victim cell. As both are the processes of victim cells acquiring or losing electrons, their likelihoods are affected by the data written to the victim cells [12], [58].

RowHammer and RowPress have different access patterns and bitflip characteristics [15], [25], [29], [39]. RowHammer, which has an attack pattern of repeatedly activating and precharging a single row, can unintentionally affect cell values in its physically-adjacent rows. RowPress, which keeps a single row activated for a long time, can cause errors in its physically nearby rows with a much lower activation count (the number of ACT-PRE command pairs applied to the row in the intervals of its adjacent rows being refreshed). Unlike RowHammer, which causes bitflips regardless of the cell's charge, RowPress specifically induces bitflips only in the charged state [39]. Accordingly, the factors affecting the AIB phenomenon can be categorized into four specific types: 1) the attack patterns (RowHammer vs. RowPress), and 2) the type of gates (the neighboring gate vs. the passing gate).

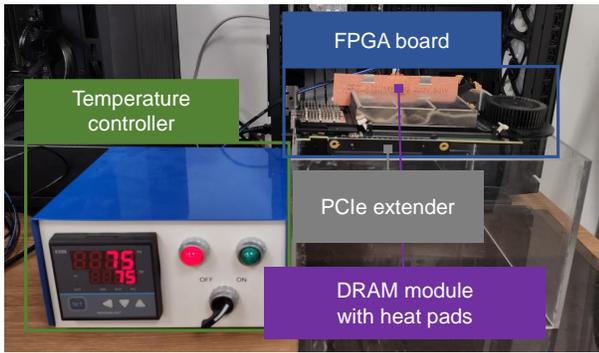


Figure 4. Test infrastructure for DDR4 DIMM and HBM2.

### III. EXPERIMENTAL METHODOLOGY

We first introduce the experimental setup and three utilized reverse-engineering techniques: AIB, RowCopy, and retention-time test. We point out common pitfalls in reverse-engineering procedures that often stem from complexities in physical to DRAM address mapping [70] and twisted data pin connection to each chip. Such pitfalls can be avoided based on publicly available information, which is yet scattered across various documents.

#### A. FPGA-based Testing Infrastructure

We modified SoftMC [14] and DRAM Bender [47] to execute the three DRAM reverse-engineering techniques and setup an FPGA-based DRAM testing platform (Figure 4). We constructed our FPGA-based DRAM testing platform using Xilinx Alveo U200 [78] and U280 [79] for testing DDR4 and HBM2, respectively. We tested 376 DDR4 chips from three major DRAM manufacturers (160 chips from Mfr. A, 128 chips from Mfr. B, and 88 chips from Mfr. C), and 4 HBM2 stacks from Mfr. A [80] for our experiments (*cf.* Table I for more details). Additionally, we presume no aging effect in DRAM devices’ fault rate according to the previous works [42], [65], [68]. Therefore, we conducted our experiments on various DRAM chips from 2016 to 2021 using the same experimental methodology. We controlled the DRAM testing platform to issue consecutive DRAM commands to DDR4 and HBM2 with a minimum interval of 1.25ns and 1.67ns (equal to tCK), respectively. We also employed a temperature controller and silicon rubber heaters to control the temperature of DRAM chips. We performed our experiments with DDR4 DIMMs at 75°C,<sup>3</sup> whereas we tested HBM2 at a constant room temperature as we could not regulate the temperature of HBM2.

#### B. DRAM Reverse-engineering Techniques

We utilize the three DRAM reverse-engineering techniques to uncover the DRAM microarchitecture and operations, and analyze the DRAM AIB characteristics.

<sup>3</sup>Although RowHammer [25], [29], [49], [50] and RowPress [39] are both known to exhibit temperature-dependent error behaviors, we did not observe significant differences in trends at other temperatures, which did not change our key observations and conclusions.

Table I  
THE TESTED DDR4 AND HBM2 CHIPS.

DRAM type	Vendor	Chip type	Density	Year	# chips
DDR4	Mfr. A	×4	8Gb	2016	80
		×4	8Gb	2017	16
		×4	8Gb	2018	32
		×4	8Gb	2021	32
		×8	8Gb	2017	16
		×8	8Gb	2018	32
DDR4	Mfr. B	×8	8Gb	2019	64
		×8	8Gb	2017	32
		×8	8Gb	2018	24
		×8	8Gb	2019	8
DDR4	Mfr. C	×4	8Gb	2018	32
		×4	8Gb	2021	32
		×8	8Gb	2016	8
HBM2	Mfr. A	4-Hi stack	4GB/stack	N/A	4
		4-Hi stack	4GB/stack	N/A	4

**Activate-induced bitflips (AIBs)** can indicate which row is adjacent to the activated aggressor row based on the fact that the physically most adjacent row is affected most [25], [73]. Most rows have two physically adjacent rows (above and below). However, the row at the edge of a subarray boundary has only one physically adjacent row. The causes of AIBs are the injection/removal of electrons into/from the victim cell and capacitive crosstalk, depending on whether the WL of the aggressor row is a neighboring gate or a passing gate (§II-D).

**RowCopy** [62] is an out-of-specification in-memory operation that copies the value of one row to another row within the same subarray using charge-sharing between a BL and a cell. First, a source row is activated. After tRAS, the row is precharged. However, if the destination row is activated soon enough, the BL will not be fully precharged to  $V_{dd}/2$  yet. Because the capacitance of a BL is much larger than a cell, the source row values can be effectively copied to the destination row through a charge transfer from the BL to the cell. We identify the height of the subarrays because RowCopy is not possible between other subarrays. Also, exploiting the fact that adjacent subarrays share half of the SAs in an open bitline structure, we identify the type of subarray structure (open or folded bitline) for each tested DRAM.

**Retention time test** [36] allows us to correctly distinguish between true-cells and anti-cells. Value 1 is represented by charged and discharged states in true-cells and anti-cells, respectively. This is a design choice to reduce the noise or optimizing the data path from the SAs to the I/O [36]. The DRAM cells naturally leak charge over time, which leads to retention failure unless periodically refreshed. The retention time of a cell is the length of time before it loses its data. Exploiting the fact that leakage occurs from a charged state to a discharged state, we perform a retention time test to distinguish between true-cells and anti-cells. We discovered that only true-cells are used in Mfr. A and Mfr. B’s DRAM chips, whereas both the true-cells and anti-cells are interleaved at a subarray granularity in Mfr. C’s DRAM chips.

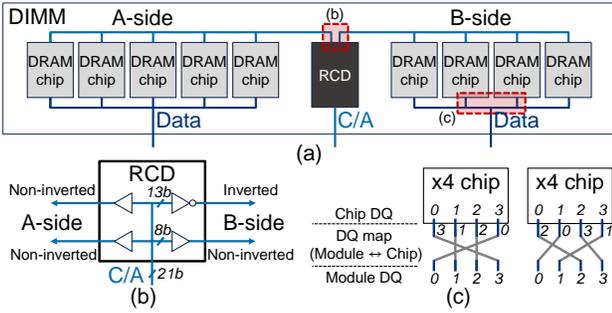


Figure 5. Identified DRAM mapping for reverse-engineering.

### C. Common Pitfalls from Address and Data Mapping

While complexities in physical-to-DRAM address mapping and data pin (DQ) twisting (or remapping) are disclosed in publicly available documents, they are regularly overlooked, leading to the following common pitfalls; (1) row address remapping at a registered clock driver (RCD) chip, (2) row address remapping in a DRAM chip, and (3) DQ remapping at each chip.

**Common pitfall-(1):** The row address can be remapped at the RCD chip [21]. The RCD chip of the registered DIMM (RDIMM) or load-reduced DIMM (LRDIMM) reduces the MC’s driving load to broadcast command/address (C/A) signals to multiple DRAM chips by decoupling the the C/A signals driven by the MC from all the DRAM chips (Figure 5(a)). As illustrated in Figure 5(b), *address inversion* is enabled by default for the RCD chip to conserve power and reduce simultaneous output switching current [21]. When the address inversion is turned on, some of the row and bank addresses to the DRAM chips on B-side (right in Figure 5(b)) are inverted, while A-side (left in Figure 5(b)) receives the non-inverted address. Such inversion can be easily neglected, which can lead to misinterpreted observations, such as direct non-adjacent RowHammer effect<sup>4</sup> [25], half-row [85], and incorrectly interpreted spare rows [3], [66]. When the inversion was considered, we were not able to observe the three phenomena (concurring with prior study [4]), whereas we could reproduce them when we disregarded such inversion. We took the inversion fully into account in our analysis, similar to prior studies [4].

**Common pitfall-(2):** The row address can also be remapped by the internal remapping scheme of each DRAM chip. For example, while DRAM decoders may preserve the sequential row order when mapping from physical addresses to DRAM row addresses, they can also scramble the row order. In a similar way to prior studies [25], [29], [50], we reconstructed internal row mapping by executing single-sided RowHammer attacks. The two rows with the most errors are the physically most adjacent rows. We found that only DDR4 and HBM2 of Mfr. A remapped rows internally, while DDR4 of Mfr. B and Mfr. C did not. From now on, we base our analysis on the remapped row addresses.

<sup>4</sup>A phenomenon where frequently activating  $N^{th}$  row can *directly* affect not only distance 1 (i.e.,  $N \pm 1^{th}$  rows) but also distance 3, 5, and further away rows.

Table II  
TERMINOLOGIES USED THROUGHOUT §IV

Symbol	Description
Data swizzling	The data reordering that occurs when transferred from MC to DRAM, and vice-versa.
$RD_{data}$	The amount of data that is read from a single chip for a single RD command (e.g., 32-bit). i.e., cache-line width divided by the number of chips.
Edge subarray	The subarray that is at the physical edge, with only one neighboring subarray.
<i>Even/odd BL</i>	BL that is indexed by an even/odd number.

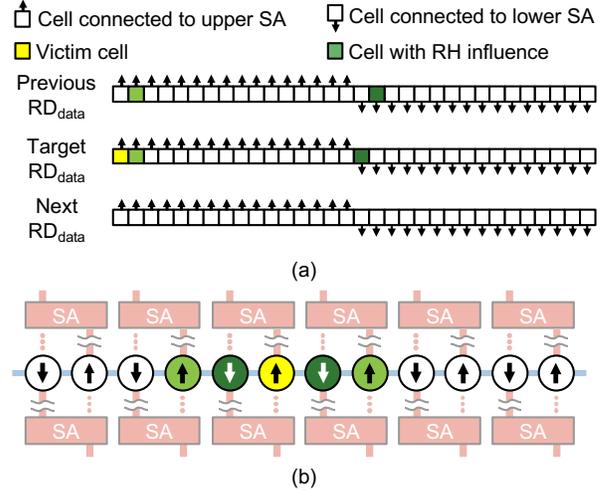


Figure 6. (a) An example of reverse-engineering the data swizzling and (b) its cell layout view with open BL structure. We identify the most influenced cells to each cell (victim cell) for RowHammer and distinguish whether they are connected to an upper SA or a lower SA by RowCopy.

**Common pitfall-(3):** Most prior studies [3], [13], [25], [29] commonly use data patterns, such as 0x55 or 0xAA assuming a straightforward connection. However, we note that DQ pins from a DIMM to each DRAM chip are also remapped (Figure 5(c)) [19], [43], [44]. Therefore, even though the user writes 0x55, each DRAM chip can receive different data (e.g., 0x33, 0xCC, or 0x99). To write the same data into all DRAM chips, we thoroughly took DQ twisting into consideration.

## IV. MACROSCOPIC ANALYSIS OF DRAM MICROARCHITECTURE

We conduct a *macroscopic* analysis on DRAM microarchitecture. First, we reverse-engineer the *data swizzling* that occurs between MC to DRAM, which is utilized across the rest of the analysis. Then, multiple observations are presented in the order of MAT, row, and subarray. Table II summarizes the terminologies that are frequently used in this section.

### A. Data Swizzling and MAT Structure

We reverse-engineered the data swizzling based on AIB and identified the MAT width. We faithfully considered DQ twisting, the horizontal AIB influence (**O11**), and the *even/odd BL* distinguished by RowCopy. We define the *even/odd BL* as the BL indexed by an even/odd number when we set an ordered number starting from the physically leftmost BL. The

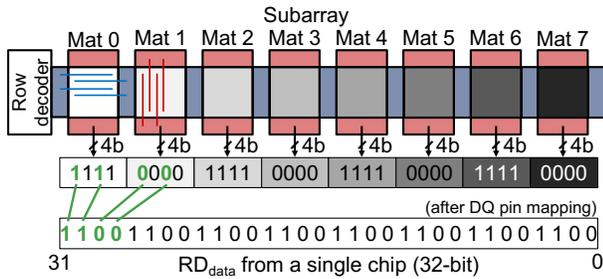


Figure 7. The specific manner in which 32-bit  $RD_{data}$  is physically stored in a Mfr. A’s DDR4  $\times 4$  chip.

horizontal influence, or the fact that physically adjacent *victim* cells affect AIB, is elaborated later (§V).

First, based on the observed fact that horizontally adjacent cells impact AIB, we find the set of cells that are adjacent to a certain victim cell, in a brute-force way. Figure 6(a) illustrates our testing methodology. For each specific victim cell, we identified four different victim cells with the largest influence. Some were in the same  $RD_{data}$  while others were in the adjacent  $RD_{data}$ . We define the  $RD_{data}$  as the data being read from a single *chip* for a single RD command (e.g., 32-bit for a  $\times 4$  chip). For example, from Mfr. A DDR4  $\times 4$  chips, we observed that bit 0 of an  $RD_{data}$  is influenced by bit 1 and 16 of the same  $RD_{data}$ , and bit 1 and 17 of the previous  $RD_{data}$ . Repeated experiments granted us the set of horizontally adjacent cells. However, because the distance  $\pm 1$  and  $\pm 2$  cells have indistinguishable differences in influence (O11), full mapping could not be acquired.

Second, we utilized RowCopy to distinguish the *odd* and *even BL*, allowing us to distinguish distance  $\pm 1$  and  $\pm 2$  cells. As explained before, in the open bitline structure, half of the cells (e.g., odd) are connected to the upper SAs, while the other half (e.g., even) are linked to the lower SA. Exploiting this, we discover if each of the collected adjacent cells have *even* or *odd BL* (Figure 6(a)). When we sequentially examine each cell, checking its four adjacent cells and which has odd/even BL, we can gain the full data swizzling as shown in Figure 6(b).

Lastly, while we identified cells in a row that influence each other, we also recognized a set of cells that are isolated from each other. We speculate such isolation is due to peripheral circuits between the MATs, such as the local row decoder and sub-WL drivers (see Figure 7). This makes it difficult for a cell in one MAT to affect a cell in another MAT. Thus, we group such isolated cells while sweeping every  $RD_{data}$ , which indicates the width of a MAT. In the tested  $\times 4$  DDR4 chips, the measured MAT widths are 512-bit, 1024-bit, and 512-bit for Mfr. A, Mfr. B, and Mfr. C, respectively.

Based on these processes, we could reconstruct the final data swizzling as shown in Figure 7. The 32-bit  $RD_{data}$  of Mfr. A’s  $\times 4$  DRAM chip is collected from 8 different MATs, each provisioning 4-bit. Each 4-bit from a MAT is again reorganized, as shown in the green line of the figure. However, while we numbered each MAT from 0 to 7 for convenience, we could not figure out the physical ordering of each MAT.

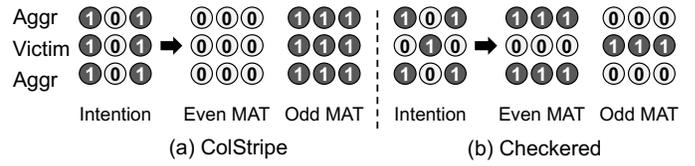


Figure 8. Unintended arrangement of data in even and odd MATs for commonly used data patterns (ColStripe and Checkered) without considering DRAM internal column address mapping.

**Observation-1:** The data of a single RD command is collected from multiple MATs and reorganized due to data swizzling.

**Observation-2:** The MAT width, or the number of cells in a row within a single MAT, is measured to be 512- or 1024-bit for tested  $\times 4$  DDR4 chips.

Our newly found data swizzling also suggests that the impact of previously understood data patterns on AIB is imprecise. For example, a ‘ColStripe’ pattern [3], [25], [50], [71], which alternates the data for every BL, actually acts as a ‘Solid’ pattern without proper mapping (Figure 8(a)). Similarly, a ‘Checkered’ pattern acts as a ‘RowStripe’ pattern (Figure 8(b)). Our detailed analysis on the data patterns in association with the newly discovered mapping and  $6F^2$  is provided in §V.

### B. Coupled-row Activation

A coupled-row activation is identified for certain  $\times 4$  DRAM chips. Both RowCopy and AIB indicate that when a row is activated (e.g.,  $i$ th row), its coupled row (e.g.,  $(i + N_{row}/2)$ th row) is activated as well. We dub such two rows indexed by two different rows specified by physical address but mapped to the same single DRAM row as a coupled-row pair. Such a behavior was exhibited by Mfr. A and Mfr. B’s  $\times 4$  DDR4 chips and Mfr. A’s HBM2. We speculate that this is a result of an optimization that ensures multiple DRAM types (e.g.,  $\times 4$  and  $\times 8$ ) to maintain the same density of cells per WL (e.g., 4096-bit or 8192-bit), regardless of the DRAM I/O width. This can serve as another AIB vulnerability, unless the host is aware of this coupling and applies proper mitigation to both the victim row and its coupled row.

**Observation-3:** For some DRAM chips, activating a row can result in the unintended activation of the coupled row.

### C. Subarray Structure

The subarray height of a DRAM chip can be accurately verified using the RowCopy-based reverse-engineering technique, and cross-checked with AIBs. Half of the cells share SAs with the upper or lower subarrays in an open bitline structure (§II-A). Thus, when we look for row address boundaries where RowCopy starts to work only for half of the cells, we can identify (1) the subarray boundary and (2) which two subarrays are adjacent. Through experiments, we have discovered that all

Table III  
THE STRUCTURE OF SUBARRAYS AND ROWS WE IDENTIFIED.

DRAM type	Vendor	Chip type	Year	Subarray composition	Edge subarray interval	Coupled-row distance
DDR4	Mfr. A	×4	2016	11 × 640-row	per	64K rows
			2017	2 × 576-row (per 8192-row)	16K rows	
		2018	2021	4 × 832-row	per	N/A
				1 × 768-row (per 4096-row)	32K rows	
		×8	2017	11 × 640-row	per	N/A
			2019	2 × 576-row (per 8192-row)	16K rows	
	Mfr. B	×4	2019	4 × 832-row (per 4096-row)	per	64K rows
			2017	2018	4 × 832-row	per
		1 × 768-row (per 4096-row)			32K rows	
		×8	2019	2 × 688-row	per	N/A
			2021	1 × 672-row (per 2048-row)	32K rows	
		Mfr. C	×4	2016	1 × 688-row	per
2018	2 × 680-row (per 2048-row)			4K rows		
×8	2016		2 × 688-row	per	N/A	
	2019		1 × 672-row (per 2048-row)	32K rows		
HBM2	Mfr. A	4-Hi	N/A	4 × 832-row 1 × 768-row (per 4096-row)	per 8K rows	8K rows

tested chips have an open bitline structure.<sup>5</sup> Also, while Mfr. A and Mfr. B copied the data in an inverted form due to the SA structure, Mfr. C copied the data as is because Mfr. C has true-/anti-cell interleaved at the subarray granularity (§III-B).

AIB can also be utilized to identify the subarray heights [4], [37], [38]. SAs that are more than 100 times larger than a DRAM cell [41], [56] separate two different subarrays, preventing AIB from occurring between two rows that are physically separated by SAs. Consequently, only the aggressor row belonging to the same subarray and adjacent to the victim row can cause bitflips. Prior studies also exploited RowHammer to find subarray heights [4], [48] or used the DRAM command sequence (ACT-PRE-ACT) to derive the adjacency between two subarrays [82]. However, because RowCopy can provide information on both aspects and is more efficient in time than the AIB method, we mainly used RowCopy and relied on AIB only for validation.

Utilizing the aforementioned methodology, we discovered that the subarray height is not a power of 2 and also can vary within a single chip, in contrast to the common conception. Mfr. A’s DDR4 ×4 chips (2016 and 2017) and ×8 chips (2017 and 2019) have a repeated pattern of 11 subarrays with 640 rows and two subarrays with 576 rows (a total of 8192 rows). The other DDR4 chips from Mfr. A have a pattern of four subarrays with 832 rows and one subarray with 768 rows (a

<sup>5</sup>Depending on the manufacturer, generation, and type, the exact bit/column location where only half of the row succeeds in RowCopy differs. This could be due to variations in the column decoder or connection between local and global I/O.

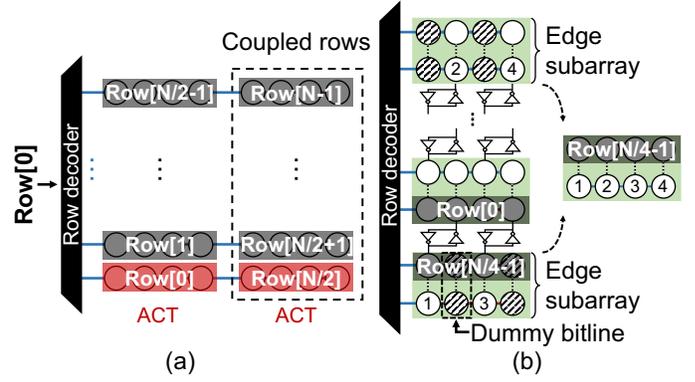


Figure 9. Characteristics of subarray structures: (a) Activating a row could incur activating its coupled row. (b) An edge subarray physically consists of a pair of subarrays, each having dummy bitlines.

total of 4096 rows) is repeated. We have also identified that the Mfr. B’s DDR4 chips and Mfr. A’s HBM2 chips have identical subarray structures with the up-to-date (till 2021) Mfr. A’s DDR4 DRAM. By contrast, Mfr. C’s DDR4 chips made in 2016 and 2018-2021 have a pattern of one subarrays with 688 rows and two subarray with 680 rows (a total of 2048 rows) and 688 rows and one subarray with 672 rows (a total of 2048 rows), respectively. We presume that the varying height of the subarray is a compromise between deteriorating timing parameters and higher cell density and fewer SAs when the cell per BL (subarray height) increases. This concurs with the trend that the subarray height has been increasing with the DRAM technology scaling. Table III summarizes the discovered subarray compositions.

**Observation-4:** The subarray heights are not power of 2, and different across different generations and within a chip.

We also observed that two different edge subarrays, or subarrays that are at the physical edge with only one neighboring subarray, work in tandem to create a single full subarray. While we found out that most subarrays are sequentially adjacent, following the row address order, we identified corner cases. For example, for some tested DRAM chips, when RowCopy was executed for the 0<sup>th</sup> row as a source and the  $(N_{row}/4 - 1)^{th}$  row as a destination, half of the bits were copied despite the large difference in the row address.  $N_{row}$  denotes the total number of rows in a bank (e.g.,  $2^{17}$ ). Because the 0<sup>th</sup> row and the  $(N_{row}/4 - 1)^{th}$  row belong to the subarrays of the bottom and top edges, respectively, we speculate that these two subarrays work together as a single subarray. We observed that most of DDR4 chips manufactured in 2018–2021 have edge subarrays at every 16K rows (×4:  $N_{row}/8$ , ×8:  $N_{row}/4$ ) or 32K rows (×4:  $N_{row}/4$ , ×8:  $N_{row}/2$ ) boundary regardless of manufacturer. However, Mfr. C’s DDR4 ×8 chips manufactured in 2016 have edge subarrays at every 4K rows (×8:  $N_{row}/16$ ) boundary. The Mfr. A’s HBM2 chip was 32K rows ( $N_{row}/4$ ) and 8K rows ( $N_{row}/2$ ).

Two subarrays working in tandem is reasonable considering the fact that only half of the cells are connected to SAs on

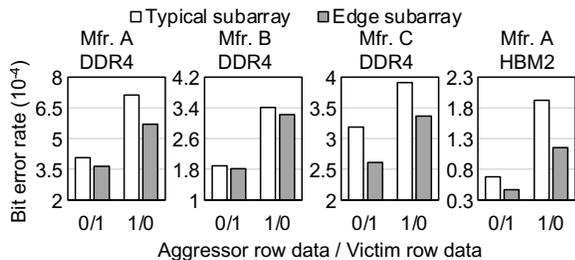


Figure 10. Comparison of AIB-induced BERs based on the subarray type (typical vs. edge subarrays) when the data pattern of an aggressor row is the inverse of a victim row.

Table IV  
SYMBOLS AND TERMINOLOGIES USED THROUGHOUT §V

Symbol	Description
$\{i, j\}$	DRAM cell with WL $i$ and BL $j$ .
Top/bottom cell	Type of cells that are isomorphic to each other, as shown in Figure 11.
<i>Even/odd WL</i>	WL that is indexed by an even/odd number.
$\text{Vic}_0$	Tested victim cell for data pattern experiment (e.g., $\{1,2\}$ ).
$\text{Vic}_{-2,-1,1,2}$	Adjacent victim cells with distance -2, -1, 1, and 2 (e.g., $\{1,0\}$ , $\{1,1\}$ , $\{1,3\}$ , and $\{1,4\}$ ).
$\text{Aggr}_{-2,-1,0,1,2}$	Adjacent aggressor cells with distance -2, -1, 0, 1, and 2 (e.g., $\{2,0\}$ , $\{2,1\}$ , $\{2,2\}$ , $\{2,3\}$ , and $\{2,4\}$ for upper aggressor row).

either side of the edge (see Figure 9). It is aligned with what the prior open bitline structures [26], [35], [69] proposed. Two edge subarrays only connect half the bitlines to the SAs, while the other half bitlines are left as dummies. Thus, when accessing an edge subarray, a simultaneous access to two rows (one for a pair of edge subarrays) is necessary to form a full single subarray.

**Observation-5:** For certain DRAM chips with the open bitline structure, two edge subarrays work in tandem to create a single full subarray.

Lastly, we also observed that the edge subarrays demonstrate a uniquely lower bit error rate (BER) for AIB (see Figure 10). We tested two different data patterns and measured the AIB-induced BERs of all the subarrays: for (aggressor, victim), (0,1) and (1,0). We identified that for both DDR4 and HBM2, edge subarrays exhibited lower BER, especially when the aggressor data was 1. We attribute such a difference to the dummy BL of the edge subarrays. Because only half the BLs are used in the edge subarray, the unused dummy BLs may preserve the precharge voltage state [26], [35]. This can be backed up by the fact that the BER is lower when the aggressor value is 1, which indicates that the dummy BLs are at least not full  $V_{dd}$ .

**Observation-6:** Edge subarrays exhibit lower BER from AIB, which can be attributed to dummy BLs.

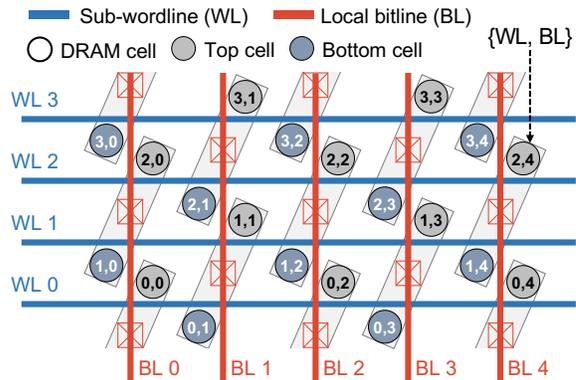


Figure 11.  $6F^2$  DRAM cell layout of  $4 \times 4$  cell array. Each cell represents its position through indices of  $\{WL, BL\}$ .

## V. MICROSCOPIC ANALYSIS OF AIB CHARACTERISTICS

In this section, we extend our analysis to *microscopic* AIB error characteristics, leveraging the  $6F^2$  cell structure and the accurate data swizzling (**O1**). Table IV summarizes the frequently used symbols and terms throughout this section.

### A. Top and Bottom Cells in $6F^2$ Structure

We adopt a coordinate notation for each cell, to be more concise in our discussion. For example,  $\{1,0\}$  denotes the cell that is connected to  $WL1$  and  $BL0$  (Figure 11). Exploiting the regular pattern of cell arrays with a  $6F^2$  structure, we categorize all the DRAM cells into two types; top and bottom cells with respect to their relative locations within a P-substrate shared by a pair of cells. Every top (or bottom) cell is isomorphic to each other. For a top cell, its upper aggressor row forms a passing gate, whereas its lower aggressor row becomes a neighboring gate. The opposite holds for a bottom cell. For instance, the passing gate of the top cell  $\{1,1\}$  is the upper  $WL2$ , whereas its neighboring gate is the lower  $WL0$ . Another noticeable pattern is that for a fixed row, the top and bottom cells appear in an alternating manner as the BL index increases (from left to right). We also define an *even/odd WL*, which refers to WL that is indexed by an even/odd number (e.g.,  $WL0$  in Figure 11 is an *even WL*).

### B. $6F^2$ -induced AIB Characteristics

For RowPress, we use an attack pattern of 8K single-sided attack with 7.8 $\mu$ s for each activation. For RowHammer, we use 300K single-sided attack<sup>6</sup> with 35ns for each activation. Based on our experiment, the gradient for flipped cells overlapping with RowPress and RowHammer converges to 0. For the tested 1024 rows, we identified a repeating trend in error with 32-bit granularity. Thus, Figure 12 reports the average BER when the bit index is modulo 32. Also, because a victim row with *odd WL* demonstrates a similar yet reversed error pattern with the *even WL* case, we only report the *even WL* case. While

<sup>6</sup>While double-sided RowHammer attacks can induce more errors with the same number of activation count, it only complicates the error characterization. We used a single-sided attack with enough activation count to compensate for the sufficient number of errors.

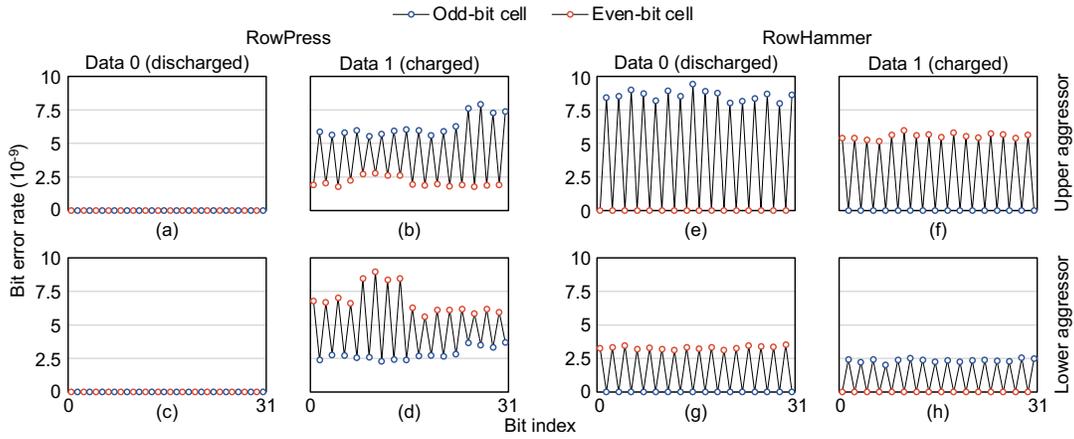


Figure 12. Average bit error rate (BER) resulting from (a)-(d) RowPress and (e)-(h) RowHammer in a Mfr. A DDR4  $\times 4$  chip. The errors occur in (b)(d)(f)(h) a charged or (a)(c)(e)(g) a discharged state, induced by (a)(b)(e)(f) an upper or (c)(d)(g)(h) a lower aggressor row. We employ the physically remapped bit index and aggregate the error across repeated 32-bit sequences, taking into account the size of  $RD_{data}$ .

the figure only summarizes the tested results of Mfr. A-2021 DDR4, similar behavior was observed in other manufacturers and HBM2. We emphasize that our analysis is highly dependent on accurate data swizzling reverse-engineering.

**RowPress:** Our RowPress experiment enabled us to successfully observe an alternating error pattern as expected. Figure 12(a,b,c,d) reports the BER rate of a certain fixed (*even WL*) victim row, for the upper or lower aggressor row. RowPress-induced bitflip was observed only in victim data 1 (charged state), as reported before [39]. In the charged state, we observed an alternating BER pattern as the bit index (BL index) increased. This accurately reflects the fact that top and bottom cells appear to take turns within a victim row (Figure 11), meaning that passing and neighboring gates appear in an alternating way, for a fixed aggressor direction (e.g., upper). Moreover, the BER pattern is *reversed* when the direction of the aggressor changes (upper and lower) or the victim row changes (*even* and *odd WL*). Such inversion can be explained by the reversed-symmetrical structure of  $6F^2$ .

**Observation-7:** RowPress occurs in an alternating pattern, which is in a reversed form between the upper/lower aggressor and *even/odd WL* victim row.

**RowHammer:** We also recognized an alternating error pattern in the RowHammer experiment. Figure 12(e,f,g,h) reports the BER rate of a particular (*even WL*) victim row for the upper and lower aggressor row. Examining the charged state victim row, we again observe a similar trend of alternating BER as the bit index increases. Also, such alternation is reversed between the upper/lower aggressor row and the *even/odd WL* victim row. Similar alternation is observed for the discharged state as well, yet in a reversed form.

**Observation-8:** RowHammer occurs in an alternating pattern, in a reversed form between the upper/lower aggressor row, *even/odd WL* victim row, and charged/discharged victim row.

Examining the gate types for each of the four tested situations (data 1/0 and upper/lower aggressor), we recognized that RowHammer can also happen on two gate types (see Figure 13). We denote the gate types as A and B because we could not fully determine whether A is a passing and B is a neighboring gate or the opposite, unless we refer to the prior study.<sup>7</sup> Moreover, we discovered that, against RowHammer, each cell is only affected by one type of gate type at a time, which is reversed when written data changes. For example, the cell with a bit index 0 demonstrates susceptibility against the upper aggressor row for data 1, and against the lower aggressor row when data is 0. Considering that the gate types on the upper and lower are the opposite for both the top and bottom cells, we can conclude that (1) RowHammer occurs in both gate types and (2) the susceptible gate type is reversed when written data changes.

**Observation-9:** RowHammer occurs at both the neighboring and passing gates.

**Observation-10:** Against RowHammer, a victim cell is only susceptible to one gate type (passing/neighboring) at a time, which is reversed when written data changes (charged/discharged).

### C. AIB Data Pattern Dependence

Based on our reverse-engineering of data swizzling and  $6F^2$ , we execute a detailed study on the data pattern dependence on AIB. Also, we propose an adversarial data pattern of victim/aggressor rows given a single victim cell based on the analysis. While prior studies [22], [25], [29], [30], [50] considered several types of data patterns, they were not based

<sup>7</sup>Prior works [12], [58] claimed that the failure mechanisms of RowHammer are charge injection from either neighboring gate (NG) or passing gate (PG), when victim cell's data is 1 or 0, respectively. However, considering the relationship between charge state of the victim cell and gate type, the characteristics of RowPress are opposite to those of RowHammer, which is different from the previous study [15]. Therefore, it is difficult to determine the gate type.

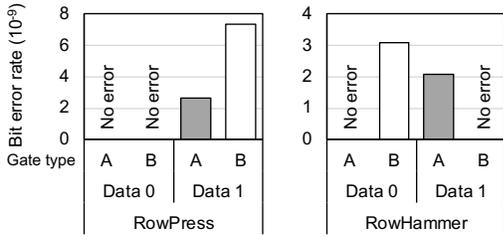


Figure 13. BER resulting from RowPress and RowHammer. The errors occur in a charged state or a discharged state, induced by a gate type A or B. A/B gates can be either passing/neighbor or neighboring/passing gates.

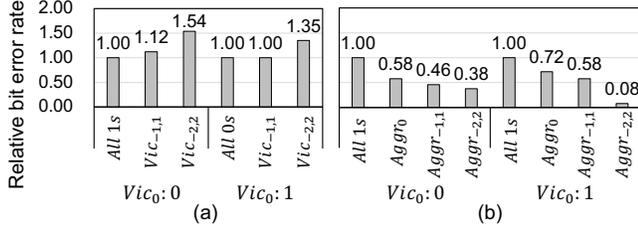


Figure 14. Relative BER influenced by (a) victim cells at distance-one and distance-two and (b) aggressor cells at distance-zero, distance-one, and distance-two, ordered in a descending manner.

upon an accurate data swizzling. This resulted in inaccurate data pattern mapping (§III-C), resulting in misinterpretations. **Victim row horizontal data pattern:** We observed that for a particular victim cell of  $Vic_0$  (e.g.,  $\{1,2\}$  in Figure 11), its adjacent four victim cells of  $Vic_{-2,-1,1,2}$  (e.g.,  $\{1,0\}$ ,  $\{1,1\}$ ,  $\{1,3\}$ , and  $\{1,4\}$ , respectively) affect the RowHammer vulnerability. Figure 14(a) summarizes the result, whose baseline BER is when all cells of aggressor and victim rows are in either (zeros, ones) or (ones, zeros). When testing the impact of four adjacent victim cells, we did not alter the data of the aggressor row. First, we recognized that BER is the worst when all four adjacent victim cells  $Vic_{-2,-1,1,2}$  hold the opposite value of the  $Vic_0$ . Also, the impact of distance-two cells  $Vic_{-2,2}$  is more significant than the impact of distance-one adjacent victim cells  $Vic_{-1,1}$ . For example, when cell  $Vic_0$  was 0, altering the value of cells  $Vic_{-1,1}$  resulted in  $1.12\times$  BER, whereas changing  $Vic_{-2,2}$  caused  $1.54\times$  BER. When  $Vic_0$  was 1, the increase in BER was  $1.00\times$  and  $1.35\times$ , respectively. We explain this to be attributed to the  $6F^2$  structure, where the *physical* distance difference between  $Vic_{-1,1}$  and  $Vic_{-2,2}$  is smaller than  $2\times$ , and physical distance between  $Vic_{-1,1}$  is also not identical (see Figure 11).

**Observation-11:** Given a particular victim cell ( $Vic_0$ ), its horizontally adjacent four victim cells ( $Vic_{-2,-1,1,2}$ ) affect its RowHammer vulnerability, which is the strongest in  $Vic_{-2,2}$ .

**Aggressor row horizontal data pattern:** We observed a horizontal data pattern dependence on RowHammer on the aggressor row. Given a particular victim cell ( $Vic_0$ ), we denote the directly adjacent aggressor cell as  $Aggr_0$  (e.g.,  $\{2,2\}$ ) and its four adjacent aggressor cells as  $Aggr_{-2,-1,1,2}$  (e.g.,  $\{2,0\}$ ,  $\{2,1\}$ ,  $\{2,3\}$ , and  $\{2,4\}$ ). While previously recognized data pattern dependence was mainly limited to  $Aggr_0$ , we newly discovered that  $Aggr_{-2,-1,1,2}$  impact the RowHammer bitflip.

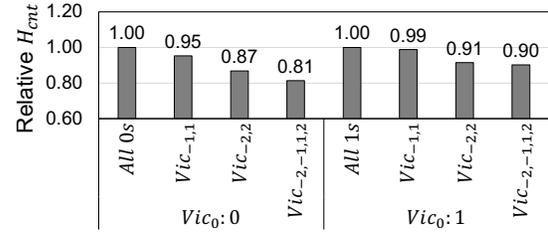


Figure 15. Relative  $H_{cnt}$  as data changes in the other victim cells.  $Vic_n$  means that data different from  $Vic_0$  was written to a cell  $n$  away from  $Vic_0$ . The aggressor row is all 1s (0s) when  $Vic_0 = 0$  (1).

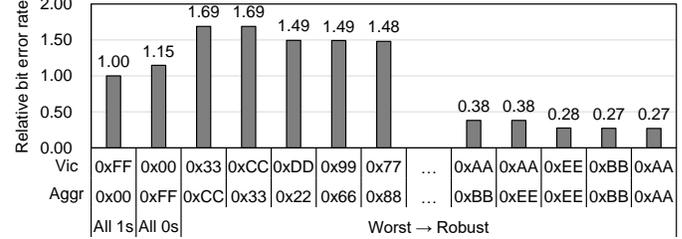


Figure 16. BER according to the data pattern of the aggressor and victim rows. Both the aggressor and victim row exhibit a repeating 4-bit sequence pattern. We represent the data pattern with values actually written to the MAT.

In fact,  $Aggr_{-2,-1,1,2}$  are more influential than  $Vic_{-2,-1,1,2}$ . With the baseline of aggressor and victim row of (zeros, ones) and (ones, zeros), we measured the BER of  $Vic_0$  while only varying the value of  $Aggr_{-2,-1,1,2}$  (Figure 14(b)). Changing  $Aggr_0$  value decreased BER by  $0.58\times$  ( $0.72\times$ ) when  $Vic_0$  was 0 (1). Altering  $Aggr_{-1,1}$  and  $Aggr_{-2,2}$  resulted in  $0.46\times$  ( $0.58\times$ ) and  $0.38\times$  ( $0.08\times$ ) drop in BER for  $Vic_0$  of 0 (1). Unlike the victim cell cases, the influence of horizontally adjacent aggressor cells was the largest when it was closest to  $Vic_0$ .

**Observation-12:** Given a particular victim cell ( $Vic_0$ ), not only its closest aggressor cell ( $Aggr_0$ ) but also horizontally adjacent aggressors ( $Aggr_{-2,-1,1,2}$ ) impact the RowHammer bitflip, which is the strongest when physically closest.

#### D. Adversarial Data Pattern for $H_{cnt}$ and BER

Leveraging the accurate data swizzling and horizontal data pattern dependence, we introduce an adversarial data pattern in terms of minimum activation count to cause the first bitflip ( $H_{cnt}$ ) and BER of the whole victim row.

**Adversarial data pattern for  $H_{cnt}$ :** We identified that setting  $Vic_{-2,-1,1,2}$  and  $Aggr_{-2,-1,0,1,2}$  as the opposite value of  $Vic_0$  deteriorates the  $H_{cnt}$  value by up to  $0.81\times$ . Due to the definition of  $H_{cnt}$ , the adversary can only target a particular victim cell,  $Vic_0$ , instead of the whole victim row. Compared to the baseline  $H_{cnt}$  where victim and aggressor row values are both 0s (1s), setting the  $Vic_{-1,1}$ ,  $Vic_{-2,2}$ , and  $Vic_{-2,-1,1,2}$  to the opposite value of 1s (0s) decreased  $H_{cnt}$  by  $0.95\times$  ( $0.91\times$ ),  $0.87\times$  ( $0.91\times$ ), and  $0.81\times$  ( $0.90\times$ ), respectively (see Figure 15). The fact that  $H_{cnt}$  decreases by a larger amount for  $Vic_{-2,2}$  than  $Vic_{-1,1}$  concurs with prior observation (O11).

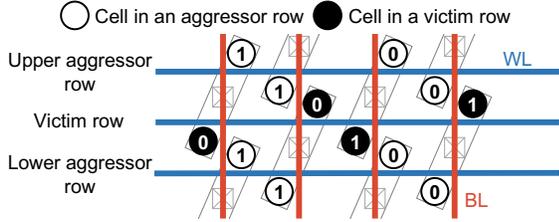


Figure 17. Among the data patterns where the data in the victim rows and the aggressor rows repeat 4-bit each, the worst-case data pattern can cause the largest number of RowHammer-induced bitflips in the victim row. The data of the aggressor row and the victim row are opposite, each data repeating two 0s or 1s. The data of the same bit index in the victim row and the aggressor rows are reversed, and two 0s and two 1s are repeated in each row.

**Observation-13:**  $H_{\text{cnt}}$  is lowered by up to  $0.81\times$ , when all cells of  $\text{Vic}_{-2,-1,1,2}$  and  $\text{Aggr}_{-2,-1,0,1,2}$  hold the opposite value of  $\text{Vic}_0$ .

**Adversarial data pattern for BER:** We identified an adversarial aggressor and victim row data pattern that deteriorates the overall BER of the victim row by up to  $1.69\times$ . We sweep 16 different data arrangements that have repeated 4-bit patterns for both victim and aggressor, testing 256 combinations in total. Figure 16 summarizes the result. The baseline is the BER when the victim and aggressor row have  $0x\text{FF}$  and  $0x00$  patterns, respectively. Among the tested 256 combinations, the worst case was when the victim and aggressor rows have  $0x33$  and  $0xCC$ , resulting in  $1.69\times$  higher BER than the baseline (Figure 17). Noticeably, this adversarial pattern is when the vertically adjacent aggressor and victim cells (e.g.,  $\text{Vic}_i$  and  $\text{Aggr}_i$ ) hold the opposite value, with a repeating pattern of 2-bits. The reason why the two-bit repeating pattern is worse than the one-bit alternating pattern can be explained by (O11), or that  $\text{Vic}_{-2,2}$  is more influential than  $\text{Vic}_{-1,1}$ .

**Observation-14:** The overall BER from AIB-induced bitflips can deteriorate by up to  $1.69\times$  when the victim and aggressor rows hold data pattern that repeats  $0x33$  and  $0xCC$ .

## VI. NEW VULNERABILITIES AND PROTECTIONS

In this section, we first examine the impact of our findings on existing AIB attacks and defenses. We propose a simple yet effective AIB protection solution against our findings. We also investigate the security implications and use cases of our observations at both macroscopic and microscopic levels.

### A. Exacerbating Known AIB Attacks

**Coupled-row activation:** Coupled-row activation (O3) can circumvent existing AIB defense mechanisms by breaking their assumption that bitflips only occur in the rows adjacent to the tracked aggressor rows. We can envision the following scenario. Suppose the two coupled rows are denoted as row-A and row-B. When an attacker only activates row-A, the AIB protection scheme is likely to only track row-A, especially when the solution is located at an MC. If the protection solution is victim row refresh-based [53], it can still be secure

by unintentionally refreshing victims of row-B. However, when the protection solution is based on state-of-the-art MC-side row swapping [60], [75], it can be neutralized because it will only relocate row-A, without affecting row-B.

Coupled-row activation also deceives the existing activation counter structures. When an attacker splits DRAM activations into each of a coupled-row pair respectively (e.g., row-A and row-B being coupled), the counter structures perceive attacker’s activations as two different row activations. However, two rows in a coupled-row pair compose a single DRAM row; thus, an attacker can easily bypass the counter structure by splitting the activations into the two rows in a coupled-row pair. Even if the MC-side AIB protection scheme acquires the coupled-row information, although itself challenging [61], the area or performance cost might become prohibitive. Most AIB mitigation schemes adopt SRAM- or CAM-based row tracking structures [27], [29], [32], [53]. Coupled-row activation doubles the number of rows to track (performance cost from twice the mitigative actions) or effectively doubles the number of activations (area cost from a larger tracking table).

Coupled-row activation aids the AIB attackers from the system perspective. Most AIB attacks utilize memory templating/messaging techniques [7], [30], [57], [72] to prepare the AIB (hammering) phase. In the memory messaging/templating phase, an attacker controls the system memory space to locate the target (victim) memory page to conduct AIB (hammering). Coupled-row activation increases the probability of successful memory messaging as the attacker concurrently accesses a single attacker-controlled memory page and benign pages belonging to other processes. Accordingly, coupled-row activation ensures a higher probability of guaranteeing adjacency between the attacker and victim pages, posing a severe threat.

**AIB adversarial pattern:** The AIB adversarial data pattern we propose worsens the existing AIB attacks 1) by decreasing the effective activation counts for an attack and 2) providing complex data pattern dependence of AIBs. Also, existing data pattern-aware AIB attacks [22], [30] should be modified to successful attacks. Pinpoint RowHammer [22] and RAMbleed [30] assume AIBs are only affected by row-wise (vertical) data patterns. However, our findings suggest that the influence of the column-wise (horizontal) data pattern should be considered. Throughout this, it is possible to increase the accuracy of the existing data pattern-aware AIB attacks.

### B. Revising Existing AIB Protection

**Protecting coupled-row activation:** To prevent bitflips caused by coupled-row activation, existing AIB defense mechanisms must consider a coupled row for each activation. Based on Table III, the relationship between coupled rows can be expressed using a simple calculation (e.g.,  $(n, n + 64K)$  for  $n : \{n \in \mathbb{N} | 0 \leq n < 64K\}$ ). Therefore, the existing tracking-based mitigation should consider coupled rows by additionally tracking its coupled row for each activation. Furthermore, considering the varying relationships with coupled rows per DRAM, in-DRAM AIB mitigation will be promising to handle the coupled-row activation.

The recently introduced Directed Refresh Management (DRFM) command is one of the prominent solutions for protecting coupled-row activation. The DDR5 JEDEC standard introduced DRFM commands as an AIB mitigation [20]. DRFM works as follows: An MC samples the activated DRAM row on `PRE` commands, and DRAM stores the corresponding DRAM row address. Then, when a DRAM receives the DRFM command, DRAM refreshes the physically adjacent rows with the stored row address. There have been attempts to utilize an RFM interface to mitigate AIB attacks [27], [40], [74]. Existing RFM-based AIB mitigation methods track DRAM activation commands and efficiently send an RFM command to trigger in-DRAM AIB mitigation actions. If the DRAM manufacturers disclose the coupled-row relationship information in either the DRAM chip’s mode register or the DRAM module’s Serial Presence Detect chip, an MC can read the information using Mode Register Read commands or `i2c` interface. Then, an MC can effectively track both coupled-row activations as a single aggressor row’s activation. Thus, coupled-row-related AIB attacks can be mitigated by adopting a DRFM-based AIB mitigation with a minimum modification.

**Protecting adversarial data pattern AIB attacks:** Data scrambling [24], [45], [83] can be an efficient mitigation technique to defend against adversarial data pattern AIB attacks. Data scrambling obfuscates the data written to the memory devices by masking the original data with the randomly generated bitmask or using encryption algorithms. Originally, Intel proposed data memory scrambling to enhance the resistance to irregular signals and power noise. Recently, it has been revealed that data scrambling can be used for security enhancements [6], [24], [83]. Modern processors manufactured by two major processor vendors (*e.g.*, Intel and AMD) enable the MC-side memory scrambling or encryption by default [24], [45], [83]. The adversarial data pattern we propose consists of both row-wise and column-wise data pattern (**O11-14**). Therefore, through a more robust PRNG algorithm involving both row and column addresses for generating bitmasks, the memory scrambling method can guarantee strong AIB adversarial data pattern resistance as in the case of the cold boot attacks [83]. Lastly, adversarial data pattern-aware error correcting code (ECC) algorithm/design and coding theory could be promising mitigation techniques [8], [28], [59], [64], [76].

### C. Possible New Attacks and Use Cases of Our Findings

Our findings on DRAM microarchitectures and operations introduce new threats to DRAM-based memory systems. The presence of edge subarrays and coupled-row activation leads to differences in power consumption based on the DRAM and subarray type. The activation of edge subarrays triggers two activations of rows in each different subarray, doubling the DRAM dynamic power consumption. Similarly, coupled-row activation doubles the power consumption due to the activation occurring in an arbitrary row. If DRAM power consumption is analyzed, it is possible to distinguish which memory (row and subarray) is accessed. Thus, an additional analysis of DRAM power-based side-/covert-channel could be intriguing [2].

Also, our findings enhance the reliability and robustness of processing in memory (PIM) techniques. For example, well-known PIM techniques, such as in-memory row copy operations (RowCopy) [10], [62] or bitwise in-memory operation using many row activation [10], [11], pose a significant threat to computer systems. For example, considering coupled-row activation, RowCopy operations in some  $\times 4$  DRAM chips or HBM2 cause unauthorized data copy. Unauthorized data copy of unintended DRAM rows compromises the confidentiality of modern computer systems. Therefore, processing using memory [62] necessitates a deeper understanding of precise DRAM operations. We believe that future DRAM-based memory system research must consider DRAM microarchitectures, associated operations, and characteristics for security.

## VII. RELATED WORK

There has been a large body of DRAM experimental analysis and characterization research, such as analyzing retention time variation [36] and latency variation [31] in commercial DRAM chips. After the advent of DRAM read disturbance error (RowHammer), a number of works have attempted to analyze and characterize the DRAM AIB characteristics [3], [13], [25]. Cojocar et al. [3] provide the method to evaluate the AIB vulnerability in the cloud environments. Kim et al. [25] demonstrate an experimental characterization of AIB on real DRAM chips and evaluate AIB mitigation techniques based on their characterization results. Hassan et al. [13] uncover undocumented in-DRAM TRR mechanisms on real DRAM chips. Also, there are works that experimentally analyze the effects of various factors (such as temperature [2], [49], [50] and wordline voltage [81]) on AIB characteristics. A new type of AIB attack called RowPress [39] was recently introduced. However, to the best of our knowledge, DRAMScope, which extended [46], is the first work that deeply considers DRAM  $6F^2$  cell structure and exact DRAM internal mappings on DRAM AIB characterization.

## VIII. CONCLUSION

We have reliably revealed the DRAM microarchitectures, associated behaviors, and activate-induced bitflip (AIBs) characteristics through AIB tests, retention time tests, and RowCopy using commercial DRAM chips. We showed that precise mapping information of DRAM modules and chips is necessary to accurately analyze the AIB characteristics. We discovered undisclosed DRAM microarchitectures and associated behaviors, such as dummy bitline, edge subarray, and coupled-row activation. We clarified the common misconceptions from prior DRAM studies, such as the non-adjacent AIB phenomenon and fixed height of subarrays. By considering the DRAM’s microscopic aspect, such as DRAM  $6F^2$  cell structure, we also identified the data pattern dependency on the AIB phenomenon. We anticipate that our new observations, clarifications, and the experimental methodology will enrich future DRAM AIB experimental analyses, and AIB attacks and defenses.

## ACKNOWLEDGEMENTS

This work was supported by Samsung Electronics Co., Ltd (IO201207-07812-01), an Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2020-0-01300, No. 2021-0-01343, and IITP-2023-RS-2023-00256081), and a grant from PRISM, one of the seven centers in JUMP 2.0, a Semiconductor Research Corporation (SRC) program sponsored by DARPA. Nam Sung Kim has a financial interest in Samsung Electronics. The EDA tool was supported by the IC Design Education Center (IDEC), Korea. The ICT at Seoul National University (SNU) provides research facilities. Jung Ho Ahn, the corresponding author, is with the Department of Intelligence and Information and the Interdisciplinary Program in Artificial Intelligence, SNU.

## APPENDIX

### ARTIFACT APPENDIX

#### A. Abstract

Our artifact provides guidelines, source code, and scripts for reproducing the figures in the paper. We offer FPGA-based infrastructures for experiments, including modified SoftMC [14] and DRAM Bender [47]. The experiments of DRAMScope consist of 1) RowHammer attack, 2) RowCopy operation, 3) Retention time test, and 4) RowPress attack for DDR4 and HBM2. We provide scripts to analyze the experimental results and plot the figures presented in this paper.

#### B. Artifact check-list (meta-information)

- **Program:** FPGA-based infrastructure: Modified SoftMC [14] and DRAM Bender [47] for Table IV and Figure 10, 12, 13, 16.
- **Compilation:** C++14 for FPGA-based infrastructure and Python 3.6.9 for analysis.
- **Metrics:** Address remapping and bit error rate.
- **Output:** CSV files are generated from FPGA experiments. *PNG* and *SVG* files are graphs, similar to the figures in this paper, generated by scripts.
- **How much disk space required (approximately)?:** 2.2GB for an FPGA-based infrastructure and 1GB for the results of DRAMScope.
- **How much time is needed to prepare workflow (approximately)?:** 1 hour.
- **How much time is needed to complete experiments (approximately)?:** 5 hours.
- **Publicly available?:** Yes.
- **Code licenses (if publicly available)?:** The MIT License for prior works (*i.e.*, SoftMC [14] and DRAM Bender [47]) and DRAMScope.
- **Archived?:** Yes. <https://zenodo.org/records/11044630>

#### C. Description

1) *How to Access:* Our modified FPGA-based infrastructure, script files, and instructions of our experiments are publicly available at GitHub repository ([https://github.com/scale-snu/AE\\_DRAMScope\\_ISCA2024](https://github.com/scale-snu/AE_DRAMScope_ISCA2024)) and Zenodo (<https://zenodo.org/records/11044630>). Prior works for our FPGA-based infrastructure are publicly available at:

- **SoftMC:** <https://github.com/CMU-SAFARI/SoftMC>

- **DRAM Bender:** <https://github.com/CMU-SAFARI/DRAM-Bender>

2) *Hardware Dependencies:* We utilize the following FPGA-based infrastructure:

- A host x86 machine supporting PCIe 3.0  $\times$ 16 slots
- An FPGA board with DIMM slots supported by DRAM Bender [47] (e.g., Xilinx Alveo U200 [78] and U280 [79])
- An FPGA board with HBM2 supported by modified SoftMC [14] for HBM2 (e.g., Xilinx Alveo U280 [79])
- A rubber heater attached to the DIMM for temperature control
- A temperature controller connected to the rubber heater

We conducted experiments using Intel Core i5-7500 (Kaby Lake) CPU and i5-8400 (Coffee Lake) CPU. We experimented with RDIMMs containing x4 and x8 chips, utilizing RDIMMs from Samsung (e.g., *M393A2K40BB1-CRC*), SK Hynix (e.g., *HMA84GR7JJR4N-WM*), and Micron (e.g., *MTA18ASF2G72PZ-2G9*).

3) *Software Dependencies:*

- GNU Make 4.1+
- C++14 build toolchain
- Python 3.6.9+
- AMD Vivado 2020.2+
- pip packages: *matplotlib* and *seaborn*
- Ubuntu 18.04 (Linux kernel 5.4.0-150-generic)

#### D. Installation and Experiment Workflow

For more detailed guidelines, please refer to `README.md` files in the following repository: [https://github.com/scale-snu/AE\\_DRAMScope\\_ISCA2024](https://github.com/scale-snu/AE_DRAMScope_ISCA2024)

#### E. Evaluation and Expected Results

After conducting each experiment on an FPGA-based infrastructure, the results files, including addresses where bitflips occurred or other experiment results, are automatically generated. We use scripts to generate the figures in this paper based on the data in the result files. Please refer to the `README.md` files in the repository for the detailed processes.

## REFERENCES

- [1] M. V. Beigi, Y. Cao, S. Gurumurthi, C. Recchia, A. Walton, and V. Sridharan, "A Systematic Study of DDR4 DRAM Faults in the Field," in *HPCA*, 2023.
- [2] Y. Cohen, K. S. Tharayil, A. Haenel, D. Genkin, A. D. Keromytis, Y. Oren, and Y. Yarom, "HammerScope: Observing DRAM Power Consumption Using Rowhammer," in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022.
- [3] L. Cojocar, J. Kim, M. Patel, L. Tsai, S. Saroiu, A. Wolman, and O. Mutlu, "Are We Susceptible to Rowhammer? An End-to-End Methodology for Cloud Providers," in *2020 IEEE Symposium on Security and Privacy (SP)*, 2020.
- [4] L. Cojocar, K. Loughlin, S. Saroiu, B. Kasikci, and A. Wolman, "mFIT: A Bump-in-the-Wire Tool for Plug-and-Play Analysis of Rowhammer Susceptibility Factors," *Technical Report-Microsoft Research*, 2021.
- [5] L. Cojocar, K. Razavi, C. Giuffrida, and H. Bos, "Exploiting Correcting Codes: On the Effectiveness of ECC Memory Against Rowhammer Attacks," in *2019 IEEE Symposium on Security and Privacy (SP)*, 2019.
- [6] V. Costan and S. Devadas, "Intel SGX Explained," *Cryptology ePrint Archive*, 2016.

- [7] F. de Ridder, P. Frigo, E. Vannacci, H. Bos, C. Giuffrida, and K. Razavi, "SMASH: Synchronized Many-sided Rowhammer Attacks from JavaScript," in *30th USENIX Security Symposium*, 2021.
- [8] A. Fakhzadehgan, Y. N. Patt, P. J. Nair, and M. K. Qureshi, "SafeGuard: Reducing the Security Risk from Row-Hammer via Low-Cost Integrity Protection," in *HPCA*, 2022.
- [9] P. Frigo, E. Vannacc, H. Hassan, V. Van Der Veen, O. Mutlu, C. Giuffrida, H. Bos, and K. Razavi, "TRRespass: Exploiting the Many Sides of Target Row Refresh," in *2020 IEEE Symposium on Security and Privacy (SP)*, 2020.
- [10] F. Gao, G. Tziantzioulis, and D. Wentzloff, "ComputeDRAM: In-Memory Compute Using Off-the-Shelf DRAMs," in *MICRO*, 2019, p. 100–113.
- [11] F. Gao, G. Tziantzioulis, and D. Wentzloff, "FracDRAM: Fractional Values in Off-the-Shelf DRAM," in *MICRO*, 2022.
- [12] S. K. Gautam, S. K. Manhas, A. Kumar, and M. Pakala, "Mitigating the Passing World Line Induced Soft Errors in Saddle Fin DRAM," *IEEE Transactions on Electron Devices*, vol. 67, no. 4, pp. 1902–1905, 2020.
- [13] H. Hassan, Y. C. Tugrul, J. S. Kim, V. Van der Veen, K. Razavi, and O. Mutlu, "Uncovering In-DRAM RowHammer Protection Mechanisms: A New Methodology, Custom RowHammer Patterns, and Implications," in *MICRO*, 2021.
- [14] H. Hassan, N. Vijaykumar, S. Khan, S. Ghose, K. Chang, G. Pekhimenko, D. Lee, O. Ergin, and O. Mutlu, "SoftMC: A Flexible and Practical Open-Source Infrastructure for Enabling Experimental DRAM Studies," in *HPCA*, 2017.
- [15] S. Hong, D. Kim, J. Lee, R. Oh, C. Yoo, S. Hwang, and J. Lee, "DSAC: Low-Cost Rowhammer Mitigation Using In-DRAM Stochastic and Approximate Counting Algorithm," *arXiv preprint arXiv:2302.03591*, 2023.
- [16] B. Jacob, D. Wang, and S. Ng, *Memory systems: cache, DRAM, disk*. Morgan Kaufmann, 2010.
- [17] JEDEC, "High Bandwidth Memory DRAM(HBM1,HBM2)," 2015.
- [18] JEDEC, "DDR4 SDRAM Standard," 2017.
- [19] JEDEC, "DDR4 SDRAM Registered DIMM Design Specification," 2019.
- [20] JEDEC, "DDR5 SDRAM Standard," 2022.
- [21] JEDEC, "DDR4 Registering Clock Driver (DDR4RCD02)," 2023.
- [22] S. Ji, Y. Ko, S. Oh, and J. Kim, "Pinpoint Rowhammer: Suppressing Unwanted Bit Flips on Rowhammer Attacks," in *Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security*, 2019.
- [23] J. Jung and M. Erez, "Predicting Future-System Reliability with a Component-Level DRAM Fault Model," in *MICRO*, 2023, pp. 944–956.
- [24] D. Kaplan, J. Powell, and T. Woller, "AMD Memory Encryption," *White paper*, 2016.
- [25] J. S. Kim, M. Patel, A. G. Yağlıkçı, H. Hassan, R. Azizi, L. Orosa, and O. Mutlu, "Revisiting RowHammer: An Experimental Analysis of Modern DRAM Devices and Mitigation Techniques," in *ISCA*, 2020.
- [26] J.-Y. Kim, M.-G. Kang, J.-Y. Lee, and B.-H. Cho, "Semiconductor Memory Device Having Dummy Bit Line," 2015, uS Patent 9,147,440.
- [27] M. J. Kim, J. Park, Y. Park, W. Doh, N. Kim, T. J. Ham, J. W. Lee, and J. Ahn, "Mithril: Cooperative Row Hammer Protection on Commodity DRAM Leveraging Managed Refresh," in *HPCA*, 2022, pp. 1156–1169.
- [28] M. J. Kim, M. Wi, J. Park, S. Ko, J. Choi, H. Nam, N. S. Kim, J. Ahn, and E. Lee, "How to Kill the Second Bird with One ECC: The Pursuit of Row Hammer Resilient DRAM," in *MICRO*, 2023.
- [29] Y. Kim, R. Daly, J. Kim, C. Fallin, J. H. Lee, D. Lee, C. Wilkerson, K. Lai, and O. Mutlu, "Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors," in *ISCA*, 2014.
- [30] A. Kwong, D. Genkin, D. Gruss, and Y. Yarom, "RAMBleed: Reading Bits in Memory Without Accessing Them," in *2020 IEEE Symposium on Security and Privacy (SP)*, 2020.
- [31] D. Lee, S. Khan, L. Subramanian, S. Ghose, R. Ausavarungrinun, G. Pekhimenko, V. Seshadri, and O. Mutlu, "Design-Induced Latency Variation in Modern DRAM Chips: Characterization, Analysis, and Latency Reduction Mechanisms," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 2017.
- [32] E. Lee, I. Kang, S. Lee, G. E. Suh, and J. Ahn, "TWiCe: Preventing Row-hammering by Exploiting Time Window Counters," in *ISCA*, 2019, pp. 385–396.
- [33] S. Lee, K. Kim, S. Oh, J. Park, G. Hong, D. Ka, K. Hwang, J. Park, K. Kang, J. Kim, J. Jeon, N. Kim, Y. Kwon, K. Vladimir, W. Shin, J. Won, M. Lee, H. Joo, H. Choi, J. Lee, D. Ko, Y. Jun, K. Cho, I. Kim, C. Song, C. Jeong, D. Kwon, J. Jang, I. Park, J. Chun, and J. Cho, "A 1.25V 8Gb, 16Gb/s/pin GDDR6-based Accelerator-in-Memory supporting 1TFLOPS MAC Operation and Various Activation Functions for Deep-Learning Applications," in *2022 IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 65, 2022, pp. 1–3.
- [34] S. Lee, S.-h. Kang, J. Lee, H. Kim, E. Lee, S. Seo, H. Yoon, S. Lee, K. Lim, H. Shin, J. Kim, O. Seongil, A. Iyer, D. Wang, K. Sohn, and N. S. Kim, "Hardware Architecture and Software Stack for PIM based on Commercial DRAM Technology: Industrial Product," in *ISCA*, 2021.
- [35] K.-N. Lim, W.-J. Jang, H.-S. Won, K.-Y. Lee, H. Kim, D.-W. Kim, M.-H. Cho, S.-L. Kim, J.-H. Kang, K.-W. Park, and B.-T. Jeong, "A 1.2 V 23nm 6F2 4Gb DDR3 SDRAM with Local-bitline Sense Amplifier, Hybrid LIO Sense Amplifier and Dummy-less Array Architecture," in *ISSCC*, 2012.
- [36] J. Liu, B. Baiyen, Y. Kim, C. Wilkerson, and O. Mutlu, "An Experimental Study of Data Retention Behavior in Modern DRAM Devices: Implications for Retention Time Profiling Mechanisms," in *ISCA*, 2013.
- [37] K. Loughlin, J. Rosenblum, S. Saroiu, A. Wolman, D. Skarlatos, and B. Kasicki, "Siloz: Leveraging DRAM Isolation Domains to Prevent Inter-VM Rowhammer," in *Proceedings of the 29th Symposium on Operating Systems Principles*, 2023.
- [38] K. Loughlin, S. Saroiu, A. Wolman, and B. Kasicki, "Stop! Hammer Time: Rethinking Our Approach to Rowhammer Mitigations," in *Proceedings of the Workshop on Hot Topics in Operating Systems*, 2021.
- [39] H. Luo, A. Olgun, A. G. Yağlıkçı, Y. C. Tuğrul, S. Rhyner, M. B. Cavlak, J. Lindegger, M. Sadrosadati, and O. Mutlu, "RowPress: Amplifying Read Disturbance in Modern DRAM Chips," in *ISCA*, 2023.
- [40] M. Marazzi, P. Jatke, F. Solt, and K. Razavi, "ProTRR: Principled yet Optimal In-DRAM Target Row Refresh," in *2022 IEEE Symposium on Security and Privacy (SP)*, 2022.
- [41] M. Marazzi, T. Sachsenweger, F. Solt, P. Zeng, K. Takashi, M. Yarema, and K. Razavi, "HiFi-DRAM: Enabling High-fidelity DRAM Research by Uncovering Sense Amplifiers with IC Imaging," in *ISCA*, Jul. 2024.
- [42] J. Meza, Q. Wu, S. Kumar, and O. Mutlu, "Revisiting Memory Errors in Large-Scale Production Data Centers: Analysis and Modeling of New Trends from the Field," in *2015 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, 2015, pp. 415–426.
- [43] Micron, "DDR4 SDRAM RDIMM MTA36ASF4G72PZ – 32GB," 2014.
- [44] Micron, "DDR4 SDRAM RDIMM MTA18ASF2G72PZ – 16GB," 2015.
- [45] C. P. Mozak, "Suppressing Power Supply Noise Using Data Scrambling in Double Data Rate Memory Systems," 2011, US Patent 7,945,050 B2.
- [46] H. Nam, S. Baek, M. Wi, M. J. Kim, J. Park, C. Song, N. S. Kim, and J. H. Ahn, "X-ray: Discovering DRAM Internal Structure and Error Characteristics by Issuing Memory Commands," *IEEE Comput. Archit. Lett.*, vol. 22, no. 2, pp. 89–92, 2023.
- [47] A. Olgun, H. Hassan, A. G. Yağlıkçı, Y. C. Tuğrul, L. Orosa, H. Luo, M. Patel, O. Ergin, and O. Mutlu, "DRAM Bender: An Extensible and Versatile FPGA-based Infrastructure to Easily Test State-of-the-art DRAM Chips," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2023.
- [48] A. Olgun, M. Osseiran, A. G. Yaglikci, Y. C. Tugrul, H. Luo, S. Rhyner, B. Salami, J. G. Luna, and O. Mutlu, "Understanding Read Disturbance in High Bandwidth Memory: An Experimental Analysis of Real HBM2 DRAM Chips," *arXiv preprint arXiv:2310.14665*, 2023.
- [49] L. Orosa, U. Rührmair, A. G. Yaglikci, H. Luo, A. Olgun, P. Jatke, M. Patel, J. Kim, K. Razavi, and O. Mutlu, "SpyHammer: Using RowHammer to Remotely Spy on Temperature," *arXiv preprint arXiv:2210.04084*, 2023.
- [50] L. Orosa, A. G. Yaglikci, H. Luo, A. Olgun, J. Park, H. Hassan, M. Patel, J. S. Kim, and O. Mutlu, "A Deeper Look into RowHammer's Sensitivities: Experimental Analysis of Real DRAM Chips and Implications on Future Attacks and Defenses," in *MICRO*, 2021.
- [51] J. Park, B. Kim, S. Yun, E. Lee, M. Rhu, and J. Ahn, "TRiM: Enhancing Processor-Memory Interfaces with Scalable Tensor Reduction in Memory," in *MICRO*, 2021.
- [52] S.-W. Park, S.-J. Hong, J.-W. Kim, J.-G. Jeong, K.-D. Yoo, S.-C. Moon, H.-C. Sohn, N.-J. Kwak, Y.-S. Cho, S.-J. Baek, H.-S. Park, H. G. Yoon, B.-H. Lee, J.-S. Kim, S.-H. Hwang, L.-H. Lee, H.-J. Cho, S. Y. Cho, C.-O. Chung, K.-O. Kim, M.-S. Yoo, S.-A. Jang, S.-D. Lee, and S.-W. Chung, "Highly Scalable Saddle-Fin (S-Fin) Transistor for Sub-50nm DRAM Technology," in *2006 Symposium on VLSI Technology*, 2006. *Digest of Technical Papers.*, 2006, pp. 32–33.

- [53] Y. Park, W. Kwon, E. Lee, T. J. Ham, J. Ahn, and J. W. Lee, "Graphene: Strong yet Lightweight Row Hammer Protection," in *MICRO*, 2020, pp. 1–13.
- [54] M. Patel, G. F. de Oliveira, and O. Mutlu, "HARP: Practically and Effectively Identifying Uncorrectable Errors in Memory Chips That Use On-Die Error-Correcting Codes," in *MICRO*, 2021.
- [55] M. Patel, J. S. Kim, T. Shahroodi, H. Hassan, and O. Mutlu, "Bit-Exact ECC Recovery (BEER): Determining DRAM On-Die ECC Functions by Exploiting DRAM Data Retention Characteristics," in *MICRO*, 2020.
- [56] Rambus, "DRAM Power Model," 2010. [Online]. Available: <https://www.rambus.com/energy>
- [57] K. Razavi, B. Gras, E. Bosman, B. Preneel, C. Giuffrida, and H. Bos, "Flip Feng Shui: Hammering a Needle in the Software Stack," in *25th USENIX Security Symposium (USENIX Security 16)*, 2016.
- [58] S.-W. Ryu, K. Min, J. Shin, H. Kwon, D. Nam, T. Oh, T.-S. Jang, M. Yoo, Y. Kim, and S. Hong, "Overcoming the Reliability Limitation in the Ultimately Scaled DRAM Using Silicon Migration Technique by Hydrogen Annealing," in *IEEE International Electron Devices Meeting (IEDM)*, 2017, pp. 21–6.
- [59] G. Saileshwar, P. J. Nair, P. Ramrakhiani, W. Elsasser, and M. K. Qureshi, "SYNERGY: Rethinking Secure-Memory Design for Error-Correcting Memories," in *HPCA*, 2018.
- [60] G. Saileshwar, B. Wang, M. Qureshi, and P. J. Nair, "Randomized Row-Swap: Mitigating Row Hammer by Breaking Spatial Correlation between Aggressor and Victim Rows," in *ASPLOS*, 2022.
- [61] S. Saroiu, A. Wolman, and L. Cojocar, "The Price of Secrecy: How Hiding Internal DRAM Topologies Hurts Rowhammer Defenses," in *2022 IEEE International Reliability Physics Symposium (IRPS)*, 2022.
- [62] V. Seshadri, Y. Kim, C. Fallin, D. Lee, R. Ausavarungnirun, G. Pekhimenko, Y. Luo, O. Mutlu, P. B. Gibbons, M. A. Kozuch, and T. C. Mowry, "RowClone: Fast and Energy-Efficient in-DRAM Bulk Data Copy and Initialization," in *MICRO*, 2013.
- [63] V. Seshadri, D. Lee, T. Mullins, H. Hassan, A. Boroumand, J. Kim, M. A. Kozuch, O. Mutlu, P. B. Gibbons, and T. C. Mowry, "Ambit: In-Memory Accelerator for Bulk Bitwise Operations Using Commodity DRAM Technology," in *MICRO*, 2017, pp. 273–287.
- [64] S. M. Seyedzadeh, D. Kline Jr, A. K. Jones, and R. Melhem, "Mitigating Bitline Crosstalk Noise in DRAM Memories," in *Proceedings of the International Symposium on Memory Systems*, 2017.
- [65] T. Siddiqua, V. Sridharan, S. E. Raasch, N. DeBardleben, K. B. Ferreira, S. Levy, E. Baseman, and Q. Guan, "Lifetime Memory Reliability Data from the Field," in *2017 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT)*. IEEE, 2017, pp. 1–6.
- [66] Y. H. Son, S. Lee, S. O. S. Kwon, N. S. Kim, and J. Ahn, "CiDRA: A Cache-inspired DRAM Resilience Architecture," in *HPCA*, 2015.
- [67] V. Sridharan, N. DeBardleben, S. Blanchard, K. B. Ferreira, J. Stearley, J. Shalf, and S. Gurumurthi, "Memory Errors in Modern Systems: The Good, The Bad, and The Ugly," in *ASPLOS*, 2015.
- [68] V. Sridharan, J. Stearley, N. DeBardleben, S. Blanchard, and S. Gurumurthi, "Feng Shui of supercomputer memory positional effects in DRAM and SRAM faults," in *SC '13: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, 2013, pp. 1–11.
- [69] T. Takahashi, T. Sekiguchi, R. Takemura, S. Narui, H. Fujisawa, S. Miyatake, M. Morino, K. Arai, S. Yamada, S. Shukuri, M. Nakamura, Y. Tadaki, K. Kajigaya, K. Kimura, and K. Itoh, "A Multigigabit DRAM Technology With  $6F^2$  Open-Bitline Cell, Distributed Overdriven Sensing, and Stacked-Flash Fuse," *IEEE Journal of Solid-State Circuits*, vol. 36, no. 11, pp. 1721–1727, 2001.
- [70] A. N. Udipi, N. Muralimanohar, N. Chatterjee, R. Balasubramonian, A. Davis, and N. P. Jouppi, "Rethinking DRAM Design and Organization for Energy-constrained Multi-cores," in *ISCA*, 2010, p. 175–186.
- [71] A. J. Van De Goor and I. Schanstra, "Address and Data Scrambling: Causes and Impact on Memory Tests," in *Proceedings First IEEE International Workshop on Electronic Design, Test and Applications*, 2002.
- [72] V. Van Der Veen, Y. Fratantonio, M. Lindorfer, D. Gruss, C. Maurice, G. Vigna, H. Bos, K. Razavi, and C. Giuffrida, "Drammer: Deterministic Rowhammer Attacks on Mobile Platforms," in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016.
- [73] A. J. Walker, S. Lee, and D. Beery, "On DRAM Rowhammer and the Physics of Insecurity," *IEEE Transactions on Electron Devices*, vol. 68, no. 4, pp. 1400–1410, 2021.
- [74] M. Wi, J. Park, S. Ko, M. J. Kim, N. S. Kim, E. Lee, and J. Ahn, "SHADOW: Preventing Row Hammer in DRAM with Intra-Subarray Row Shuffling," in *HPCA*, 2023.
- [75] J. Woo, G. Saileshwar, and P. J. Nair, "Scalable and Secure Row-Swap: Efficient and Safe Row Hammer Mitigation in Memory Systems," in *HPCA*, 2023.
- [76] S. C. Woo, W. Elsassner, M. Hamburg, E. Linstadt, M. R. Miller, T. Song, and J. Tringali, "RAMPART: RowHammer Mitigation and Repair for Server Memory Systems," in *9th International Symposium on Memory Systems*, 2023.
- [77] X.-C. Wu, T. Sherwood, F. T. Chong, and Y. Li, "Protecting Page Tables from RowHammer Attacks using Monotonic Pointers in DRAM TrueCells," in *ASPLOS*, 2019.
- [78] A. Xilinx, "Alveo U200 FPGA Board." [Online]. Available: <https://www.xilinx.com/products/boards-and-kits/alveo/u200.html>
- [79] A. Xilinx, "Alveo U280 FPGA Board." [Online]. Available: <https://www.xilinx.com/products/boards-and-kits/alveo/u280.html>
- [80] A. Xilinx, "Supercharge Your AI and Database Applications with Xilinx's HBM-Enabled UltraScale+ Devices Featuring Samsung HBM2," 2019.
- [81] A. G. Yağlıkçı, H. Luo, G. F. De Oliviera, A. Olgun, M. Patel, J. Park, H. Hassan, J. S. Kim, L. Orosa, and O. Mutlu, "Understanding RowHammer Under Reduced Wordline Voltage: An Experimental Study Using Real DRAM Devices," in *2022 52nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, 2022.
- [82] A. G. Yağlıkçı, A. Olgun, M. Patel, H. Luo, H. Hassan, L. Orosa, O. Ergin, and O. Mutlu, "HiRA: Hidden Row Activation for Reducing Refresh Latency of Off-the-Shelf DRAM Chips," in *MICRO*, 2022.
- [83] S. F. Yitbarek, M. T. Aga, R. Das, and T. Austin, "Cold Boot Attacks are Still Hot: Security Analysis of Memory Scramblers in Modern Processors," in *HPCA*, 2017.
- [84] M. S. Yoo, K. S. Choi, W. K. Sun, S. G. Choi, J. I. Kim, Y. I. Son, H.-G. Choi, T. K. Oh, Y. T. Hwang, Y. Chun, J. G. Jeong, S. K. Park, J. H. Choi, S. J. Hong, and S. Park, "Saddle-fin cell transistors with oxide etch rate control by using tilted ion implantation (TIS-Fin) for sub-50-nm DRAMs," Feb 2010.
- [85] T. Zhang, K. Chen, C. Xu, G. Sun, T. Wang, and Y. Xie, "Half-DRAM: A High-bandwidth and Low-power DRAM Architecture from the Rethinking of Fine-grained Activation," in *ISCA*, 2014, pp. 349–360.
- [86] L. Zhou, J. Li, Z. Qiao, P. Ren, Z. Sun, J. Wang, B. Wu, Z. Ji, R. Wang, K. Cao, and R. Huang, "Double-sided Row Hammer Effect in Sub-20 nm DRAM: Physical Mechanism, Key Features and Mitigation," in *2023 IEEE International Reliability Physics Symposium (IRPS)*, 2023.